




Configuring and Installing IBM BladeCenter

You can deploy VMware ESX Server 2.1 on a variety of hardware, including blade servers (blades). Because of the unique hardware of blade servers, we have prepared this paper to help you maximize your ESX Server experience on IBM blade servers.



This note contains the following topics:

- [Using Blade Servers with ESX Server on page 1](#)
- [Using Blades with ESX Server and VirtualCenter on page 2](#)
- [Blade Server Hardware Requirements on page 2](#)
- [Configuring ESX Server on IBM Blade Servers on page 3](#)
- [Installing ESX Server on IBM BladeCenter on page 5](#)
- [Installing ESX Server on the First Blade from the CD-ROM on page 5](#)
- [Installing ESX Server on Additional Blades on page 6](#)
- [Performing a Remote Network Installation of ESX Server with RDM on page 6](#)
- [Post-Installation Considerations on IBM Blade Servers on page 9](#)
 - [Using NIC Teaming on IBM Blade Servers on page 9](#)
 - [Creating VLANs with NIC Teaming on IBM Blade Servers on page 10](#)
 - [Creating a VLAN with NIC Teaming for the Service Console on page 10](#)
 - [Creating VLANs with NIC Teaming for Virtual Machines on page 10](#)
- [Best Practices on IBM Blade Servers on page 10](#)
 - [Mounting USB CD-ROM Devices on page 10](#)
 - [Determining the Floppy Drive on a Blade Server on page 11](#)


Using Blade Servers with ESX Server

The advantages of blade servers include greater versatility, ease of deployment and serviceability, and cost savings. When combined with VMware ESX Server, and optionally with VMware VirtualCenter, blade servers offer you even more benefits.



In particular, using VMware products with blade servers provides:

- Improved utilization by increasing server and application density. You can consolidate applications and infrastructure services onto fewer blade servers.
- Increased operational flexibility.
- Increased resource management and use of each virtual machine.



VMware products improve the existing benefits of blade servers. By using ESX Server, you can install multiple virtual servers on a single blade, thereby containing costs and maximizing the potential of your hardware.

For more information on the benefits of VMware ESX Server, see www.vmware.com/products/server/esx_features.html.

Using Blades with ESX Server and VirtualCenter

VMware VirtualCenter is a separately purchasable product that provides virtual infrastructure management software. You have a central point of control through a single management console window for your data center's virtual computing resources. VirtualCenter manages your virtual machines as a single, logical pool of processing, networking, memory, and storage resources, thereby enabling you to manage workload and optimize resource utilization.

VirtualCenter also enables the VMotion add-on module that enables zero-downtime maintenance and ensures 100% service availability. You can migrate a running virtual machine to a different physical server, connected to the same storage area network (SAN), without service interruption. By moving virtual machines on the fly, you can perform maintenance on the underlying hardware and storage, without scheduling downtime and without impacting users.

For more information on the benefits of VMware VirtualCenter, see www.vmware.com/products/vmanage/vc_features.html.

Blade Server Hardware Requirements

Blade server hardware requirements are as follows:

- Blade server enclosure
- IBM HS20 blade servers



- Two BladeCenter 4-port Gigabit Ethernet Switch modules
- BladeCenter Fibre Channel Expansion Card (one for each blade, if you want FC connectivity)
- BladeCenter 2-port Fibre Channel Switch module (for FC connectivity)
- Sufficient physical memory to prevent virtual machine swapping from being a significant performance issue

Refer to the ESX Server documentation for more information on system requirements at www.vmware.com/support/esx21/doc/esx21admin_res.html.

Configuring ESX Server on IBM Blade Servers

Tip — Name the blades by slot number; for example, `blade1`, `blade2`, and so on.

There are two main choices for blade server storage:

- Local storage
- Storage area network (SAN) devices with Fibre Channel adapters.

Local Storage

On IBM blade servers, you need local storage (typically an IDE drive) to install the VMkernel and VMware Service Console.

ESX Server does not support “diskless” operation, where the ESX Server base installation is on a Fibre Channel SAN logical unit number (LUN), Internet SCSI (iSCSI), network-attached storage (NAS), or other external storage.

Local SCSI

Some blade systems have SCSI peripherals that can be attached to each CPU, and take up to two SCSI drives. In typical use, these SCSI drives are placed into a RAID1 (mirrored) configuration for redundancy.

On an IBM blade, the SCSI peripheral takes up one of the blade slots, thus reducing the maximum blade density by half.

SAN Storage

Fibre Channel SANs are the preferred storage media for ESX Server and VirtualCenter in a blade environment, due to the following advantages:

- This configuration doubles the blade density per blade chassis, compared with local SCSI storage on IBM blades.
- SAN storage may be shared among multiple blades (and other systems), thus allowing storage consolidation. Often, this is a much more efficient use of storage resources than dedicated, per-system, RAID-protected storage.
- IBM blade systems support redundant host bus adapters (HBAs) to meet High Availability needs.
- The storage is more reliable (RAID5 with hot spares compared to RAID1).
- Storage is unlimited compared to the storage that fits on a single local SCSI disk.
- A shared SAN is required for using VMotion with VirtualCenter.
- Images, templates, and so on, may be shared between multiple ESX Server systems.



Typical IBM BladeCenter Storage Configuration

A typical IBM BladeCenter implementation of ESX Server and VirtualCenter has a single IDE drive (20GB, 40GB, or larger) on each blade and at least several hundred gigabytes of SAN storage split into RAID5 LUNs, visible to all members of each VirtualCenter farm.

Install the VMkernel, the service console, and the virtual machine configuration (.vmtx) files on a local drive. Typically, this is the local IDE drive on each blade. However, if your BladeCenter includes a local SCSI peripheral, then we suggest you install the VMkernel, the service console, and the virtual machine configuration (.vmtx) files on the local SCSI drive.

You can then install the virtual disk (.disk) files on LUNs in your Storage Area Network devices. VMFS volumes cannot reside on an IDE drive.

VMotion

For you to use VMotion, all blades in a VirtualCenter farm need access to the same logical unit number (LUN) on a SAN. Consequently, the VMFS volumes that contain the virtual machine virtual disk (.disk) files must be on a shared SAN accessible by ESX Server.

Note: VMotion is not supported for virtual machines hosted on local storage. The virtual machine must reside on a shared SAN.

Core Dump Partition

ESX Server core dump partitions must be on a controller visible to the virtual machines (VMkernel). We recommend that you create ESX Server core dump partitions either on a local SCSI drive, or on the SAN.

We recommend that you create a separate core dump partition for each ESX server on IBM blades. For example, you can use a separate LUN for each ESX Server machine that contains both its core dump partition and the swap file, discussed in the following section.

Swap File

ESX Server swap partitions must be on a controller visible to the virtual machines (VMkernel). We recommend that you create ESX Server swap partitions on a local SCSI drive, or on the SAN.

Purchase enough physical memory to prevent virtual machine swapping from being a significant performance issue.

Depending on the number of blades and your swap usage, you may choose to allocate a dedicated LUN for swap files. Multiple swap files from multiple ESX servers can reside on this dedicated LUN. Do not store any other kind of file (virtual machine .disk files, checkpoint files, and so on) on this LUN.


Use a unique name for each blade server swap file, such as, <server_name>.vswp.

Although you can have a total of eight swap files for each ESX Server machine, you can select only one swap file through the VMware Management Interface.



Installing ESX Server on IBM BladeCenter

For a list of blade servers supported with ESX Server, see the *ESX Server Systems Compatibility Guide* at www.vmware.com/products/server/esx_specs.html.



1. You must install ESX Server on the first IBM blade by performing a standard CD-ROM-based installation. Follow the procedure described in the *ESX Server 2.1 Installation Guide*, at www.vmware.com/support/esx21/doc/esx21install_text-steps_install.html.

Note: You can use the graphical installer with IBM blades only if you are using a USB mouse plugged into a USB port.

2. There are two different procedures for installing ESX Server on subsequent, additional blades in IBM BladeCenter. Select and complete *one* of the following procedures, based on your work environment.
 - [Performing a Remote Network Installation of ESX Server with RDM on page 6](#)
 - Performing a remote network installation by using a scripted, remote installation. Follow the steps in the *ESX Server 2.1 Installation Guide* at www.vmware.com/support/esx21/doc/esx21install_script_setup_install.html.

Installing ESX Server on the First Blade from the CD-ROM

1. On the first blade server, press these two buttons:
 - a. **CD select** to associate that blade server with the CD and floppy drive.
 - b. **KVM select** to associate that blade server with the keyboard, monitor, and mouse.
2. During the ESX Server installation, the installer attempts to probe for the video and monitor setting. Since the USB mouse for IBM blades is not supported, a warning message appears. This message is expected. Choose the “Proceed with Text Installation” option when prompted.
3. After installing ESX Server from the CD-ROM, you are prompted to reboot the server. The installer attempts to eject the CD-ROM, but fails. You can then manually eject the CD-ROM by pressing the button on the CD-ROM drive.

Note: If you do not remove the CD-ROM, then the CD-ROM installation restarts once the server boots.

4. Perform the ESX Server configuration steps in the VMware Management Interface, as described in the *ESX Server 2.1 Installation Guide*. Follow the steps at www.vmware.com/support/esx21/doc/esx21install_config_install.html.



Note: We recommend you dedicate all Fibre Channel devices to the virtual machines (VMkernel).

Installing ESX Server on Additional Blades

Once you install ESX Server on a system, you can quickly deploy or provision more ESX Server systems that share the same configuration, or have a similar configuration. You can install ESX Server on additional blades either by using the ESX Server 2.1 CD-ROM (described in the preceding section) or by using a remote network installation procedure (described in the following section).

Performing a Remote Network Installation of ESX Server with RDM

You can set up an installation script that comprises the choices you want to make during the installation of ESX Server software. This script allows you to install ESX Server remotely, without having to use an ESX Server CD in the new target ESX server system.

Complete *only* one of the following:

- Perform a remote network installation with RDM.
- Perform a scripted installation.

In this paper, we describe the remote network installation with RDM. Refer to the *ESX Server 2.1 Installation Guide* at www.vmware.com/support/esx21/doc/esx21install_script_setup_install.html for the scripted installation procedure.

Note: Use the following procedure to install ESX Server on the second and any additional blades.

These are the minimum system requirements:

- Microsoft® Windows® 2000 server
- Both IBM® Director and RDM software, version 4.11 or higher, must be installed on the Windows 2000 server
- 512MB RAM minimum for running IBM Director and RDM
- DBMS, such as Microsoft SQL Server, installed on the Windows 2000 server
- DHCP server that is configured either by RDM or by an existing server with the proper pre-boot execution environment (PXE) configuration
- NFS Server with a mount point sharing the contents of the ESX Server 2.1 CD-ROM

Use RDM to boot a pre-configured ESX Server installation image. This installer then loads the necessary files from a file server.

Complete the following steps to install ESX Server using RDM:

1. Install IBM Director and RDM on a Microsoft Windows 2000 server system. During the RDM installation, you need to install all three components:
 - RDM server (integrated with IBM Director server) — This server stores all the data and the configuration information.
 - D-server — The deployment server sends files to the target blade system (the blade on which ESX Server is to be installed). An RDM server may have multiple D-servers, each



servicing a different range of IP addresses for target systems. However, one D-server is sufficient for this installation.

- Remote console (integrated with IBM Director console) — This console provides the user interface for inspecting and controlling the RDM server. Multiple IBM Director remote consoles may be connected to the RDM server, if necessary.
2. Configure the DHCP server.

A DHCP server is required for the proper operation of RDM. Refer to the RDM documentation for the procedure to configure your DHCP server.
 3. Prepare the ESX Server boot image.

After installing ESX Server on the first blade, use the VMware Management Interface to prepare an ESX Server installation floppy disk. Configure a scripted installation for DHCP operation as described in the *ESX Server 2.1 Installation Guide* at www.vmware.com/support/esx21/doc/esx21install_script_setup_install.html. When the configuration is complete, select **Download Floppy Image** and create a floppy disk image.
 4. Create a new mount point on your NFS server and export it. For more information on how to set up an NFS server, see The Linux Documentation Project HOWTO at tldp.org/HOWTO/NFS-HOWTO.
 - a. Copy the contents of the ESX Server 2.1 CD-ROM to the root of your NFS mount point. Then copy the `ks.cfg` file from the ESX Server installation floppy to the root of your NFS mount point.
 - b. Edit the `ks.cfg` file and modify the `Installation Method` line. This line should start with a `cdrom` or `url` command. Replace this line with the following:


```
nfs --server <nfsserver> --dir <nfsdir>
```

Replace `<nfsserver>` with the IP address or host name of your NFS server and replace `<nfsdir>` with the NFS mount point.
 5. Add the ESX Server boot image to the RDM server.

By default, RDM is installed in `C:\Program Files\IBM\RDM`. We assume this default location in the following steps. If you have installed RDM in a different directory, then change the directories in the following steps, accordingly.

 - a. Open the folder for `C:\Program Files\IBM\RDM\repository\environment\`.
 - b. Create a folder in that directory called `esx`.
 - c. Copy `C:\Program Files\IBM\RDM\repository\environment\etc\pxeboot.0` to `C:\Program Files\IBM\RDM\repository\environment\esx\`.
 - d. Copy `C:\Program Files\IBM\RDM\repository\environment\etc\pxeboot.cfg` to `C:\Program Files\IBM\RDM\repository\environment\esx\`.
 - e. Edit `C:\Program Files\IBM\RDM\repository\environment\etc\default` to match the following lines.

Note: Type each line on a single line, including the line starting with `APPEND`. This line appears as two lines because of the formatting in this tech note. `<nfsserver>` is the



IP address or host name of the NFS server and `<nfsmount>` is its mount point as configured in step 4.

```
DEFAULT vmlinuz
```

```
APPEND initrd=initrd.img apic ks=nfs:<nfsserver>:<nfsmount>/
ks.cfg ramdisk_size=10240
```

- f. Copy the `initrd.img` and `vmlinuz` files from the ESX Server installation floppy to `C:\Program Files\IBM\RDM\Repository\environment\esx\`.
Be sure that the file names are all lower case.
6. Find the target system (the blade on which you want to install ESX Server) in IBM Director.
 - a. Open and log into the IBM Director console program.
 - b. There are three columns in the IBM director window. Set the "group" on the left column to "physical platforms".
 - c. Find the target system in the middle column. You can identify the target system by its MAC address, IP address, machine name, or IBM machine ID. To be sure, double-click the highlighted entry representing the target system to match the MAC address. If you cannot find the target system, boot the target system by using the pre-boot execution environment (PXE) once, and the target system is boot scanned by the IBM Director server.
7. Create an installation task.
 - a. In the right "task" column of the IBM Director window, choose **Remote Deployment Manager > Custom**.
 - b. Right click on **Custom** and create a new task.
 - c. In the **Advanced** tab, there is an editable text box that contains the script that runs the task. Modify the installation task to match the following.


```
;This is command list for custom task
BOOTTYPE !LOADBOOTSTRAP environment/esx/pxeboot.0
WAKE
!!setenv
!!SHUTDOWN
END
```
8. Create a job for the target system using the custom installation by dragging the target system in the second column onto the new task that you just created. Choose **run system** in the pop-up window, and then select **execute now**.
9. Start the ESX Server installation.
If the target system is set to wake up on LAN or boot from PXE, it should find the RDM server and load the boot image `pxeboot.0` from the D-Server. The boot image then loads the `vmlinuz` and `initrd.img` files and starts the ESX Server installer. The ESX Server installer downloads files from the NFS server and continues with the installation.
10. Perform the ESX Server configuration steps in the VMware Management Interface, as described in the *ESX Server 2.1 Installation Guide* at www.vmware.com/support/esx21/doc/esx21install_config_install.html.

Note: We recommend you dedicate all Fibre Channel devices to the virtual machines (VMkernel).



- Repeat this process for each additional blade server in the BladeCenter.

Post-Installation Considerations on IBM Blade Servers

We discuss two topics in this section:

- [Using NIC Teaming on IBM Blade Servers on page 9](#)
- [Creating VLANs with NIC Teaming on IBM Blade Servers on page 10](#)

Using NIC Teaming on IBM Blade Servers

IBM blade servers have two physical network interface cards (NICs). NIC teaming (IEEE 802.3ad) is a feature of ESX Server that allows you to create a bonded NIC that spans multiple physical NICs. Each bond acts as a virtual switch that provides multiple uplinks for your use. For more information on bonds and virtual switches, see the *ESX Server 2.1 Administration Guide* at www.vmware.com/support/esx21/doc/esx21admin_virtualswitches.html.

Because the IBM HS20 blades have only two NICs per blade, a standard NIC teaming configuration dedicates both NICs to the virtual machines, leaving no NICs for the service console. To address this limitation, ESX Server allows you to dedicate both NICs to the virtual machine and create a bond, then give network access to the service console through the `vmxnet_console` module. This action results in the service console being on the same local area network (LAN) segment as the virtual machines.

- Install ESX Server as described previously.
- As root, log into the service console. Run the `vmkpcidivvy` program in the interactive mode.


```
vmkpcidivvy -i
```

 - As you have already installed ESX Server, accept all the defaults until you get to the NICs.
 - Share the NIC that was originally assigned to the VMware Service Console. Look for the NIC that is labeled with a `[c]` and change this to shared mode by typing `s`. Leave the remaining NIC assigned to the VMkernel `[v]`.
- Create the bond. Using `vi` or a similar text editor, edit `/etc/vmware/hwconfig` and add the following two lines to the end of the file, where `x` is a bond number from 0 to 9.

```
nicteam.vmnic0.team = "bond<x>"
nicteam.vmnic1.team = "bond<x>"
```

- Use the `vmxnet_console` module to give network access to the service console. Using `vi` or a similar text editor, edit `/etc/rc.local` and add the following lines, where `x` is the bond number you selected in step 3.

```
#vmxnet_console through bond<x>
/etc/rc.d/init.d/network stop
rmmod vmxnet_console
insmod vmxnet_console devName=bond<x>
/etc/init.d/network start
mount -a
```

- Reboot ESX Server for your changes to take effect unless you are creating a VLAN for the service console. See [Creating a VLAN with NIC Teaming for the Service Console on page 10](#).



Creating VLANs with NIC Teaming on IBM Blade Servers

System administrators use virtual LANs, or VLANs to increase performance and security. VLANs also improve manageability and network tuning by defining broadcast domains without the constraint of physical location.

You can create VLANs for the service console and the virtual machines.

Creating a VLAN with NIC Teaming for the Service Console

1. Create a bond between the two NICs as described in the previous section, [Using NIC Teaming on IBM Blade Servers on page 9](#). However, do not reboot ESX Server at this time.
2. Using `vi` or a similar text editor, edit `/etc/rc.local` and add the following lines to the end of the file, where `x` is the bond number you selected and `y` is the VLAN tag number (from 1 to 4095). These lines allow you to use the `vmxnet_console` module to give network access through the bond to the service console.

```
#vmxnet_console through bond<x> VLAN<y>
/etc/rc.d/init.d/network stop
rmmod vmxnet_console
insmod vmxnet_console devName="bond<x>.<y>"
/etc/init.d/network start
mount -a
```

3. Configure the physical ports on your external switch (not the switch on BladeCenter) to be trunk ports that support the VLANs IDs that you have chosen.
4. Reboot ESX Server.

Creating VLANs with NIC Teaming for Virtual Machines

1. Create a bond between the two NICs as described in [Using NIC Teaming on IBM Blade Servers on page 9](#).
2. As root, log into the VMware Management Interface and create VLANs for your virtual machines. Refer to the online help or the *ESX Server 2.1 Administration Guide* at www.vmware.com/pdf/support/esx21 for complete instructions.

Best Practices on IBM Blade Servers

This section includes some general best practices for ESX Server on IBM blades.

Mounting USB CD-ROM Devices

To mount a USB CD-ROM driver, you must symbolically link (`ln -s`) `/dev/cdrom` to `/dev/scd0`. This symbolic link is created automatically if you installed ESX Server using the ESX Server CD-ROM. If however, you installed ESX Server using a remote or network installation, then you may need to create this symbolic link manually. To create the symbolic link, type:

```
ln -s /dev/scd0 /dev/cdrom
```

To mount the USB CD-ROM manually, type:

```
mount /dev/cdrom /mnt/cdrom
```

Note: When you switch from one blade to another, the `/dev/cdrom` link and possibly the `/mnt/cdrom` directory may be deleted. If this happens, type the following to mount the USB CD-ROM manually:

```
mount /dev/scd0 /mnt/cdrom
```

If the `/mnt/cdrom` directory is also deleted, then recreate this directory by typing:



```
mkdir /mnt/cdrom
```

Determining the Floppy Drive on a Blade Server

When an IBM blade powers on, the SCSI drivers load first (if SCSI devices are present), followed by the USB drivers (if USB devices are present). Then the Fibre Channel device drivers load, if they are allocated solely to the virtual machines (VMkernel), per our recommendation.

To determine the location of the floppy drive, you need to count the number of logical units on SCSI drives, Fibre Channel LUNs, then determine whether or not USB devices are enabled (the **CD select** and **KVM select** buttons are selected) when the blade boots.

For example, when powering on a blade, you have a single logical SCSI drive and two Fibre Channel LUNs. Therefore, the floppy drive is `/dev/sdb`. The SCSI driver loads first (there's one logical SCSI drive), then the USB drivers load (you've selected both the **CD select** and **KVM select** buttons).

Mounting Floppy Drives from the Command Line

1. Determine the `/dev/sd<x>` name for the floppy disk drive, where *x* is a letter from **a** to **z**.
2. Type the following:

```
mkdir -p /mnt/floppy
mount /dev/sd<x> /mnt/floppy
```

Mounting Floppy Drives from the VMware Remote Console

To mount a physical floppy drive on a blade server from the VMware Remote Console, you need to know the `/dev/sd<x>` name for the floppy drive, where *x* is a letter from **a** to **z**.

1. In the remote console, choose **Settings > Configuration Editor**.
2. If it is not already selected, click the **Hardware** tab.
3. Click **Floppy Drive** and select **Use floppy image** (even though you are actually using a physical floppy drive).
4. Enter the appropriate `/dev/sd<x>` device name.